

DSDP3.2 Users Manual

Steven J. Benson
Yinyu Ye

Argonne National Laboratory
Mathematics and Computer Science Division
9700 South Cass Avenue
Argonne, Illinois 60439

Department of Management Sciences
The University of Iowa
Iowa City, Iowa 52242

November 15, 2000

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing, U.S. Department of Energy, under Contract W-31-109-Eng-38.

The DSDP package uses a dual scaling algorithm to solve the positive semidefinite relaxation of combinatorial optimization problems. It solves problems of the form

$$\begin{aligned}
& \text{Minimize} && C \bullet X \\
\text{(SDP)} & && \\
& \text{Subject to} && A_i \bullet X = b_i, \quad i = 1, \dots, m, \\
& && X \in K.
\end{aligned} \tag{1}$$

The variable X must be a symmetric positive semidefinite matrix. Here $K = K_1 \oplus K_2 \oplus \dots \oplus K_p$ and K_l is the cone of $n_l \times n_l$ symmetric positive semidefinite matrices. The data matrices $C, A_i \in \Re^{n \times n}$ are symmetric. The operation $C \bullet X = \text{tr } C^T X = \sum_{jk} C_{jk} X_{jk}$ and $X(\succeq) \succ 0$ means that X is positive (semi) definite. Furthermore, we assume the matrices A_i are linearly independent, meaning that $\sum_{i=1}^m y_i A_i = 0$ implies $y_1 = \dots = y_m = 0$. Matrices X that satisfy the constraints are called feasible, while the others are called infeasible.

The dual variables are $y \in \Re^m$ and the symmetric positive semidefinite matrix S . The dual problem of (SDP) can be written as follows:

$$\begin{aligned}
& \text{Maximize} && b^T y \\
\text{(DSP)} & && \\
& \text{Subject to} && \sum_{i=1}^m y_i A_i + S = C, \quad S \in K,
\end{aligned} \tag{2}$$

where b_i and y_i is the i th elements of the vectors b and y , respectively.

Under relatively mild conditions the primal and dual optimal solution pair (X^*) and (y^*, S^*) exist, and $C \bullet X^* = b^T y^*$.

1 Dual Scaling Algorithm

Following conventional notation, let

$$\mathcal{A}X = \begin{bmatrix} A_1 \bullet X & \dots & A_m \bullet X \end{bmatrix}^T \quad \text{and} \quad \mathcal{A}^T y = \sum_{i=1}^m A_i y_i,$$

Given a dual point (y, S) such that $\mathcal{A}^T y + S - C = R$ and $S \succ 0$, and a parameter $\hat{\mu} > 0$, each iteration of the dual scaling algorithm steps toward the solution of the nonlinear system

$$\mathcal{A}X = b \quad \mathcal{A}^T y + S = C \quad \hat{\mu} S^{-1} = X. \tag{3}$$

The Newton direction satisfies the linear system

$$\mathcal{A}(\Delta X) = 0 \quad \mathcal{A}^T(\Delta y) + \Delta S = -R \quad \hat{\mu} S^{-1} \Delta S S^{-1} - \Delta X = X - \mu \hat{S}^{-1}, \tag{4}$$

and DSDP solves the Schur complement of these equations, given by

$$\begin{pmatrix} A_1 S^{-1} \bullet S^{-1} A_1 & \cdots & A_1 S^{-1} \bullet S^{-1} A_m \\ \vdots & \ddots & \vdots \\ A_m S^{-1} \bullet S^{-1} A_1 & \cdots & A_m S^{-1} \bullet S^{-1} A_m \end{pmatrix} \Delta y = \frac{1}{\hat{\mu}} b - \mathcal{A}(S^{-1}) - \mathcal{A}(S^{-1} R S^{-1}). \quad (5)$$

The dual step direction Δy is used to compute $\Delta S = -\mathcal{A}^T \Delta y - R$ and a matrix X if needed. For notational convenience, we label the left-hand matrix of (5) M . For an arbitrary feasible X ,

The algorithm selects a step size β_{k+1} such that $y^{k+1} = y^k + \beta_{k+1} \Delta y$ and $S^{k+1} = S^k + \beta_{k+1} \Delta S \succ 0$. Using the dual step direction and (4), a primal matrix X satisfying the linear constraints can be calculated using

$$X = \hat{\mu} S^{-1} + \hat{\mu} S^{-1} \Delta S S^{-1}. \quad (6)$$

DSDP does not compute X until the algorithm terminates since the step direction does not depend on X ,

Given a feasible dual starting point and appropriate choices for $\hat{\mu}$ and β , convergence results in [1] show that either the new dual point (y, S) or the new primal point X is feasible and reduces the Tanabe-Todd-Ye primal-dual potential function

$$\Psi(X, S) = \rho \ln(X \bullet S) - \ln \det X - \ln \det S$$

enough to achieve linear convergence.

A more thorough explanation of the the dual scaling algorithm used to solve these problems can be found in the papers *Solving Large-Scale Sparse Semidefinite Programs for Combinatorial Optimization*[2] and *DSDP 3: Dual Scaling Algorithm for General Positive Semidefinite Programs*. The C code that implements the algorithms can be downloaded from the Web site at <http://www.mcs.anl.gov/~benson/> or <http://dollar.biz.uiowa.edu/col/>.

2 UNIX and LINUX Installation Instructions

The compressed tar file DSDP3.2.tar.gz contains an implementation of the dual scaling algorithm for positive semidefinite programming optimization problems.

1. Create the DSDP3.2 directory structure with

```
gunzip -d DSDP3.2.tar.gz
tar -xvf DSDP3.2.tar
```

This command produces the directory DSDP3.2 and several subdirectories.

2. Change directories to DSDP3.2. Several executables have been provided. If one of these runs on your architecture, proceed to step 3. Otherwise compile the executables using

```
make install
```

This command creates libraries in each of the subdirectories and the executables for DSDP.

3. DSDP3.2 can be called from MATLAB version 5.0. Run the sample problems by starting MATLAB in the **DSDP3.2** directory and typing

```
> check;
```

Compare the output with the output in **demo.solaris**. For help using the package, type **help dsdp3**.

Run the executables by switching to directory **DSDP3.2/exec** and typing

```
dsdp3 truss1.dat-s
```

Compare the output with that in **demo.solaris**. If the output from any of the tests differs significantly from the files, please report it to the developers.

The most common cause of trouble while compiling or running DSDP concerns linking the program to LAPACK. DSDP links with the LAPACK routine **dsyev**. On most architectures, this routine is called from C programs using **dsyev_**. If this calling sequence is different on your machine, modify the file **DSDP3.2/src/solver/dsdplapack.h**. This file contains the line

```
#define dsdpdsyev dsyev_
```

which can be modified to call the appropriate routine.

If you have an optimized version of LAPACK available, you may link to it by modifying the LAPACK macro in the file **DSDP3.2/src/Makefile**. Users with MATLAB 6.0 or higher can remove the link to LAPACK, since it is already included in MATLAB 6.0.

You may also have to change the MEX macro in the directory **DSDP3.2/src/Makefile**. This macro specifies the location of the MATLAB program and the mex functionality in particular.

3 Using DSDP with MATLAB

Help using the DSDP MATLAB interface can be found by simply typing **help dsdp3** in the directory **DSDP3.2**. The command

```
> [STAT, y, X] = DSDP(A,C,b)
```

attempts to solve the positive semidefinite program by using a dual scaling algorithm. The arguments **A**, **C**, and **b**, are the same as the corresponding arguments in the SDPT3 [3] package for semidefinite programming. For a problem with p blocks and m constraints, **A** is a $p \times m$ cell array and **C** is a $p \times 1$ cell array. One block may contain LP variables, and the cells corresponding to this block should be a vector array. All other cells must contain a square, symmetric, real valued matrix. The choice of sparse or dense matrix formats belongs to the user and will affect the performance of the solver. The third argument **b** is a dense column vector of length m .

The second and third output arguments return the dual and primal solutions, respectively. The first output argument is a structure with several fields that describe the solution of the problem:

obj	an approximately optimal objective value
primal	an approximately optimal objective value if convergence to the solution was detected, infeasible if primal infeasibility is suspected, and unbounded if dual infeasibility is suspected
dual	an approximately optimal objective value if convergence to the solution was detected, infeasible if dual infeasibility is suspected, and unbounded if primal infeasibility is suspected
iterates	number of iterations used by the algorithm
dual_infeasibility	the multiple of the identity matrix added to $C - \mathcal{A}^T(y)$ in the final solution to make S positive definite
gaphist	a history of the duality gap
infhist	a history of the dual infeasibility
termcode	0 solution found, 1: primal infeasible, 2: dual infeasible

There are more ways to call the solver. The command

```
> [STAT, y, X] = DSDP(A,C,b,OPTIONS,y0)
```

specifies some options for the solver. The **OPTIONS** structure may contain any of the following parameters:

gaptol	tolerance for duality gap as a fraction of the value of the objective functions [default 1e-3]
inf_tol	tolerance for stopping because of suspicion of dual infeasibility. [default 1e-8]
step_tol	tolerance for stopping because of small steps [default 1e-2]
maxit	maximum number of iterations allowed [default 100]
printyes	1, if want to display result in each iteration, else = 0 [default 1]
scale_data	1, if want to scale the data before solving the problem, else 0 [default

1]

r0 multiple of the identity matrix added to the initial dual matrix. $S0 = C - \sum A_i y_i^0 + r0 * I$. If $r0 < 0$, a dynamic selection will be used [default -1]

max_trust_radius maximum trust region radius to be used for feasible dual iterates. Values between 1 and 10 are typical. Smaller values make the solver more robust, but larger one often accelerate convergence. A negative value indicates that the default dynamic strategy should be used [default -1].

max_infeasible_trust_radius maximum trust region radius to be used for infeasible iterates. Values equal to the dimension of the problem are typically used. A negative value indicates that the default dynamic strategy should be used. [default -1].

dual_bound a bound for the dual solution. The solver stops when a feasible dual iterate with an objective greater than this value is found [default $1e+8$].

The command

```
> [STAT, y, X] = DSDP(A,C,b,OPTIONS,y0)
```

specifies an initial dual vector $y0$. The default vector consists of positive and negative ones.

If only two output are used,

```
> [STAT,y] = DSDP()
```

returns only the solver statistics structure and an approximate dual solution y .

4 Standard Output

The following is an example output from a random problem.

Iter	Primal	Dual	Infeas.	Mu	Step	Pnrm
0	0.00000000e+00	-2.47128801e+01	0.0e+00	1.0e+01	0.00	1.00
1	-1.28537231e+01	-1.64471021e+01	0.0e+00	1.2e+00	0.08	26.94
2	-1.28537231e+01	-1.59965110e+01	0.0e+00	2.3e-01	1.00	1.35
3	-1.28537231e+01	-1.59071616e+01	0.0e+00	2.1e-01	1.00	1.07
4	-1.28537231e+01	-1.55639753e+01	0.0e+00	1.5e-01	1.00	1.22
5	-1.50630757e+01	-1.53435713e+01	0.0e+00	1.4e-01	0.08	12.59

```

6  -1.51189868e+01  -1.52669146e+01  0.0e+00  1.0e-01  0.28  3.72
7  -1.51505856e+01  -1.52461229e+01  0.0e+00  1.0e-02  1.00  0.84
8  -1.52220358e+01  -1.52350997e+01  0.0e+00  9.2e-03  0.12  8.91

```

The program will print a variety of statistics for each problem to the screen.

Iter	the current iterate number
Primal	the current estimate of the primal objective function
Dual	the current dual objective function
Infeas	the infeasibility in the current dual solution This number is the multiple of the identity matrix that has been added to the dual matrix
Mu	a central path parameter. This parameter decreases to zero as the points get closer to the solution
Step	the multiple of the dual step-direction used
Pnrm	The proximity to a point on the central path: $\ S^{.5}XS^{.5} - \hat{\mu}I\ $

5 Standalone version with SDPA files

DSDP can also be run without the MATLAB environment if the user has a problem written in sparse SDPA format. These executables have been put in the directory **DSDP3.2/exec/**. The file name should follow the executable. For example,

```
> dsdp3 truss4.dat-s
```

Other options can also be used with DSDP. These should follow the SDPA filename.

- save <filename> to save the solution into a file with a format similar to SDPA.
- y0 <filename> to specify an initial dual vector.
- maxit <iter> to stop the problem after a specified number of iterations.
- gaptol <rtol> to stop the problem when the relative duality gap is less than this number.
- trust <radius> limits the step length with a trust region radius. The default number is 3. A larger radius may accelerate the convergence of the algorithm, but a smaller radius generally makes it more robust. This radius applies only when the iterates are feasible.

`-inftrust <radius>` limits the step length with a trust region radius when the iterates are not feasible. This step size is usually larger than the previous ones so that feasibility in the problem can be found more quickly.

`-help` to print a help message

6 Applying DSDP to Graph Problems

Within the directory **DSDP3.2/exec/** is a program **maxcut** which reads a file containing a graph, generates the SDP relaxation of a maximum cut problem, and solves it. For example,

```
> maxcut Max2
```

The files that contain a graph should follow the DIMACS graph format. The first line should contain two integers. The first integer states the number of nodes in the graph, and the second integer states the number of edges. Subsequent lines have two or three entries separated by a space. The first two entries specify the two nodes that an edge connects. The optional third entry specifies the weight of the node. If no weight is specified, a weight of 1 will be assigned.

The same options that apply to reading SDPA files also apply here.

A similar program reads a graph from a file, formulates a minimum bisection problem, and solves it; for example,

```
> minbisection Max2
```

reads the graph in the file **Max2** and solves this minimum bisection problem.

7 DSDP Application Programming Interface

DSDP3.2 can also be used with a set of subroutines. Examples of using the DSDP API can be found by looking at the files **DSDP3.2/src/dsdp.c** and **DSDP3.2/src/sdpa.c**, which read data from the MATLAB environment and a SDPA file, respectively, and solve the problem by using DSDP. Other files **DSDP3.2/src/maxcut.c** and **DSDP3.2/src/minbis.c** read a graph and formulate the maximum cut and minimum bisection problems. Each of these applications includes the header file **DSDP3.2/src/dsdp3.h** and links to the library **DSDP3.2/src/dsdplib.a**. All DSDP subroutines begin with the **DSDP** prefix and return an error code. A zero code indicates success, while a nonzero indicates that an error has occurred.

7.1 Basics

To call DSDP as a subroutine, the DSDP data structure must first be created with the command

```
int DSDPCreate(int,int*, int*, DSDP *);
```

The first argument is the number of constraints in the problem, the second argument is the number of blocks in the problem, and the third argument is an array of integers specifying the dimension of each block. The length of this array must be at least as long as the number of blocks in the problem and the dimension of each block corresponds the number of rows or columns of matrices in the block. The final argument should receive the address of a DSDP structure. This routine will allocate some resources for the solver and set the pointer.

After setting the data, which will be explained in the next section, DSDP must process the data with the routine

```
int DSDPSetup(DSDP);
```

This routine allocates some additional resources for the solver and computes the eigenvalues and eigenvectors of the constraint matrices. This routine should be called before solving the problem, and should be called only once for each DSDP solver created. The command

```
int DSDPSolve(DSDP);
```

attempts to solve the problem. This routine can be called more than once. For instance, the user may try solving the problem using different initial points. Each solver created should be destroyed with the command

```
int DSDPDestroy(DSDP);
```

This routine frees the work arrays and data structures allocated by the solver.

7.2 Set Data

A positive semidefinite programming programming problem may be defined by using the following “Set” routines. The first argument in each of these routines is a pointer to the DSDP solver. To set the dual objective function, the routine

```
int DSDPSetDualObj(DSDP,int,double* );
```

accepts an array of double precision variables that contain the dual objective function b . The second argument must equal the dimension of this objective vector and match the number of constraints specified the DSDPCreate() command. DSDP will copy this data but not use the array.

The data matrices can be specified by any of the following commands. The choice of data structures is up to the user, and the performance of the problem depends upon this choice of data structures. In each of these routines, the first four arguments are a pointer to the DSDP solver, the block number, and the constraint number, and the number of rows and columns in the matrix. The blocks must be numbered consecutively, beginning with the number 0. Constraints are numbered consecutively beginning with the number 1. The primal objective matrices are specified using constraint 0. The data that is passed to the DSDP solver will be used in the solver, but not modified. The user is responsible for freeing the arrays of data it passes to DSDP after solving the problem.

To set data in a dense matrix, the routine

```
int DSDPSetDenseMat(DSDP,int, int, int, double *nnz);
```

hands DSDP a pointer to an array of length $ni \times ni$, where ni is the number of rows in the matrix. Since the matrix must be symmetric, the array may be in either row major or column major form. The routine

```
int DSDPSetSparseMat(DSDP,int, int, int, double *, int *, int *);
```

specifies a matrix in sparse form. The fifth argument is an array containing the nonzeros of the matrix. The sixth argument is an array of integers specifying the row number of each nonzero. The last argument specifies the number of nonzeros in each column. Row and column number must be numbered consecutively beginning with 0. The dimension of this integer array is one more than the number of rows and columns in the matrix. The first element in this array is 0. Subsequent elements in this array equal the previous element plus the number of nonzeros in that column. These indices indicate where in the fifth and sixth arguments to find the nonzeros for a particular column. This sparse column format, or sparse row format, given the symmetry of the matrix, is the same sparse format used by MATLAB.

A diagonal matrix can be specified with the command

```
int DSDPSetDiagMat(DSDP,int, int, int, double *);
```

The last argument must be an array equal to the rows and columns of the block. This command is especially appropriate when a block of linear programming variables is part of the problem. A diagonal matrix can also be specified in sparse format by using the command

```
int DSDPSetSpDiagMat(DSDP,int, int, int,double *, int *, int);
```

In this routine, the fifth argument is an array of nonzeros, the sixth column is an array of integers specifying the row/column number, and the seventh argument specifies the the length of the arrays in the fifth and sixth arguments which should equal the number of nonzeros in this constraint block. If every element in matrix can the same value, the command

```
int DSDPSetConstantMat(DSDP,int, int, int,double);
```

will create a data structure that will make the computations in the algorithm more efficient. Similarly, a program has a constraint where one block has a single nonzero element, which equals 1.0 and is located along the diagonal. The routine

```
int DSDPSetSingletonMat(DSDP,int, int,int,int);
```

creates an appropriate data structure. The fifth argument specifies the diagonal element where the nonzero element resides. Matrices with a 2×2 block form can be specified with the following command:

```
int DSDPSetTwoTwoMat(DSDP,int,int, int,
                    int i0, int i1, double d00, double d01, double d11);
```

The integers $i0$ and $i1$ specify the two row numbers of the nonzero block, while the next three digits define the elements of the block. The elements in the upper triangular matrix are specified in row major order. Matrices with a 3×3 block form can be specified by using

```
int DSDPSetThreeThreeMat(DSDP,int,int,int,
                        int i0,int i1,int i2,
                        double d00,double d01,double d02,
                        double d11,double d12,double d22);
```

The integers $i0$, $i1$, and $i2$ specify the three row numbers of the nonzero block. The six double precision arguments that follow define the upper triangular part of the matrix in row major form. Constraint matrix blocks that have rank one can be specified by using the command

```
int DSDPSetRank1Mat(DSDP,int, int,int,double, double *);
```

The sixth argument is an array of length equal to the number of rows and columns in the block. The constraint is the outer product of this vector with itself times a scalar, which is specified in the fifth argument.

7.3 Improving Performance

The performance of the DSDP may be improved with the proper selection of parameters and data structures. DSDP accepts multiple data structures for the data matrices, including dense, sparse, and diagonal representations. The choice of data types belongs to the user and may affect the performance. In addition, DSDP may achieve greater efficiency by placing diagonal blocks before nondiagonal blocks and smaller blocks before larger blocks.

DSDP allows the user to specify the initial starting point. The command

```
int DSDPSetY0(DSDP,int, double *);
```

specifies the initial dual vector y to be used. The second argument is an integer corresponding to the number of constraints in the problem. The length of the array in the third argument must exceed this dimension. DSDP will copy this data into its own structure. The initial dual matrix can be partially controlled by using the command

```
int DSDPSetInitialInfeas(DSDP,double);
```

This command specifies a positive number r and sets $S^0 = C - \sum A_i y_i^0 + rI$, where I is the identity matrix. If $r < 0$, a default value will of r will be chosen. If S^0 is not positive definite, the routine `DSDPSolve()` will return an error code of -100 .

The algorithm also used trust regions to determine step sizes. The size of the trust regions affects the convergence and robustness of the algorithm. By default, DSDP uses a dynamic strategy. The user may specify a maximum trust region tolerance using the commands

```
int DSDPSetMaxTrust(DSDP,double);
```

and

```
int DSDPSetMaxInfeasTrust(DSDP,double);
```

The first command sets the maximum trust region radius when the dual iterates are feasible, while the second command sets the maximum trust region radius when the iterates are infeasible. The maximum radius for feasible iterates is usually about 3, while the maximum for infeasible iterates is much higher by default. The default strategies can be specified by using a negative maximum radius. Smaller trust regions generally make the solver more robust, allowing it to solve more problems, but larger trust regions often reduce the number of iterations needed to converge. The maximum steplength is determined by dividing the trust region radius by the statistic `Pnrm`. If this statistic is consistently greater than the trust region radius, an increase in this parameter may improve performance.

7.4 Set Options

A variety of different options may be set when using DSDP. The precision of the solution can be set by using the routine

```
int DSDPSetGapTol(DSDP,double);
```

This command will terminate if the solver finds a sufficiently feasible solution such that the difference between the primal and dual objective values, normalized by the magnitude of the dual

objective value, is less than the prescribed number. A tolerance of 0.001 provides roughly three digits of accuracy, while a tolerance of $1.0e - 5$ provides roughly five digits of accuracy. The command

```
int DSDPSetInfeasTol(DSDP,double);
```

specifies how small the dual infeasibility constant r must be to be an approximate solution, and the routine

```
int DSDPSetMaxIts(DSDP,int);
```

specifies the maximum number of iterates. The routine

```
int DSDPSetDualBound(DSDP,double);
```

specifies an upper bound on the dual solution. The algorithm will terminate when it finds a point whose dual infeasibility is less than the prescribed tolerance and whose dual objective value is greater than this number. DSDP may print some information to the screen. This information can be turned off by passing a zero value into the second argument of the routine

```
int DSDPPrint(DSDP,int);
```

A nonzero value asks to print the information. The command

```
int DSDPSetScaling(DSDP,double);
```

does some partial scaling of the problem. DSDP will scale the problem by implicitly dividing the primal objective matrix by this number. An appropriate value is the magnitude of the largest element in the primal objective matrix. A potential function may also be applied to the algorithm. For feasible iterates, DSDP can set the parameters \bar{z} equal to the current primal objective value and ρ to a multiple of $n + \sqrt{n}$ and can reduce the steplength until either the dual potential function

$$\psi(y) = \rho \ln(\bar{z} - b^T y) - \ln \det S$$

has been reduced or the steplength becomes smaller than $0.1/\text{Pnrm}$. This multiple of $n + \sqrt{n}$ should be greater than 1.0 and set by using the command

```
int DSDPSetPotentialParameter(DSDP, double);
```

A nonpositive parameter, which is the default, means that the potential function will not be enforced at each step.

7.5 Viewing Solution

The dual solution vector y can be viewed by using the command

```
int DSDPGetY(DSDP,int, double *);
```

The user passes the size of the dual vector and an array of double to DSDP, which will copy the dual solution into this array. The routine

```
int DSDPGetInfeasibility(DSDP, double *);
```

returns the dual infeasibility r used in the final iteration. To view the primal solution, the user must call either the command

```
int DSDPSetDenseXMat(DSDP,int, int, double *);
```

or

```
int DSDPSetDiagXMat(DSDP,int, int, double *);
```

before solving the problem. The user should use the first command for SDP blocks and the second command for linear programming variable blocks. The first, second, and third arguments to these routines are a pointer to the DSDP solver, the block number, and the number of rows and columns in the block, respectively. In the fourth argument, the user passes an array of double precision variables long enough to store either a full dense matrix of the size designated in the third argument, or a vector whose dimension equal the number of rows or columns in the block. The command

```
int DSDPNoXMat(DSDP,int){
```

specifies a block for which the dual solution is not needed. Performance in the routine `DSDPSetup()` is generally better when the nondiagonal primal solution matrices X have been requested because it uses this array as work space instead of dynamically allocating it for every matrix.

The success of DSDP can be interpreted with the command

```
int DSDPStopReason(DSDP,int *);
```

This command sets the second argument to an integer. If this integer equals 1, DSDP found a primal and dual solution. If the integer is 2 DSDP terminated due to small steps. If the integer equals 3, DSDP asserts that the problem is primal infeasible and the dual problem is unbounded, and if the integer equals 4, DSDP asserts that the problem is dual infeasible and primal unbounded. If the integer is -5 DSDP terminated after reaching the maximum number of iterations.

Information about the convergence of the solver can be obtained with the commands

```
int DSDPGetGapHistory(DSDP, int, double *);
```

and

```
int DSDPGetInfeasHistory(DSDP, int, double *);
```

These commands retrieve the history of the duality gap and the dual infeasibility for up to 100 iterations. The user passes an array of doubles and the length of this array.

8 Performance

More information about the convergence of the algorithm and performance of the software can be found at [1] and [2].

9 Acknowledgements

We thank Cris Choi and Xiong Zhang for their help in developing this code, and Hans Mittelmann for his efforts in testing and benchmarking it.

References

- [1] S. J. Benson, Y. Ye, and X. Zhang, Solving large scale sparse semidefinite programs for combinatorial optimization. *SIAM Journal of Optimization*, 10:443–461, 2000.
- [2] S. J. Benson and Y. Ye. DSDP3: Dual Scaling Algorithm for General Positive Semidefinite Programs. Technical Report ANL/MCS-P851-1000, Argonne National Laboratory, November, 2000.
- [3] K. C. Toh, M. J. Todd, and R. H. Tütüncü. SDPT3 – A MATLAB software package for semidefinite programming, version 1.3 *Optimization Software and Methods*, 11:545-581, 1999.